

Boosting Few-Shot Visual Learning with Self-Supervision

Spyros Gidaris¹, Andrei Bursuc¹, Nikos Komodakis², Patrick Pérez¹, Matthieu Cord^{1,3}
¹valeo.ai ²LIGM, Ecole des Pont ParisTech ³Sorbonne Université

Abstract

Few-shot learning and self-supervised learning address different facets of the same problem: how to train a model with little or no labeled data. Few-shot learning aims for optimization methods and models that can learn efficiently to recognize patterns in the low data regime. Self-supervised learning focuses instead on unlabeled data and looks into it for the supervisory signal to feed high capacity deep neural networks. In this work we exploit the complementarity of these two domains and propose an approach for improving few-shot learning through self-supervision. We use self-supervision as an auxiliary task in a few-shot learning pipeline, enabling feature extractors to learn richer and more transferable visual representations while still using few annotated samples. Through self-supervision, our approach can be naturally extended towards using diverse unlabeled data from other datasets in the few-shot setting. We report consistent improvements across an array of architectures, datasets and self-supervision techniques. We provide the implementation code at <https://github.com/valeoai/BF3S>.

1. Introduction

Deep learning-based methods have achieved impressive results on various image understanding tasks, such as image classification [22, 29, 52, 54], object detection [47], or semantic segmentation [4]. However, to learn successfully these tasks, convolutional neural networks (convnets) need to access large volumes of manually labeled training data. If not, the trained convnets might suffer from poor generalization performance to the test data. In contrast, humans are perfectly capable of learning new visual concepts from only a few examples, generalizing without difficulty to new data. *Few-shot learning* [8, 9, 26, 30, 55] aims to endow artificial perception systems with a similar ability, especially with the help of modern deep learning.

Hence, the goal of few-shot visual learning is to devise recognition models that are capable of efficiently learning to recognize a set of classes despite that there are available very few training examples for them (*e.g.*, only 1 or 5 examples per class). To avoid overfitting due to data scarcity, few-shot

learning algorithms rely on transfer learning techniques and have two learning stages. During the first stage, the model is usually trained using a different set of classes, called *base classes*, which is associated with a large number of annotated training examples. The goal of this stage is to let the few-shot model acquire transferable visual analysis abilities, typically in the form of learned representations, which are mobilized in the second stage. In this next step, the model indeed learns to recognize *novel classes*, unseen during the first learning stage, using only a few training examples per class.

Few-shot learning relates to self-supervised representation learning [6, 7, 14, 31, 40, 61]. The latter is a form of unsupervised learning that trains a model on an annotation-free pretext task defined using only the visual information present in images. The purpose of this self-supervised task is to make the model learn image representations that would be useful when transferred to other image understanding tasks. For instance, in the seminal work of Doersch *et al.* [6], a network, by being trained on the self-supervised task of predicting the relative location of image patches, manages to learn image representations that are successfully transferred to the vision tasks of object recognition, object detection, and semantic segmentation. Therefore, as in few-shot learning, self-supervised methods also have two learning stages, the first that learns image representations with a pretext self-supervised task, and the second that adapts those representations to the actual task of interest. Also, both learning approaches try to limit the dependence of deep learning methods on the availability of large amounts of labeled data.

Inspired by the connection between few-shot learning and self-supervised learning, we propose to combine the two methods to improve the transfer learning abilities of few-shot models. More specifically, we add a self-supervised loss to the training loss that a few-shot model minimizes during its first learning stage (see Figure 1). Hence, we artificially augment the training task(s) that a few-shot model solves during this stage. We argue that this task augmentation forces the model to learn a more diversified set of image features, and this, in turn, improves its ability to adapt to novel classes with few training data. Moreover, since self-supervision does not require data annotations, one can include extra unlabeled data to the first learning stage. By extending the size and




Figure 1: **Combining self-supervised image rotation prediction and supervised base class recognition in the first learning stage of a few-shot system.** We train the feature extractor $F_{\theta}(\cdot)$ with both annotated (top branch) and non-annotated (bottom branch) data in a multi-task setting. We use the annotated data to train the object classifier $C(\cdot)$ with the few-shot classification loss L_{few} . For the self-supervised task, we sample images from the annotated set (and optionally from a different set of non-annotated images). Here, we generate four rotations for each input image, process them with $F_{\theta}(\cdot)$ and train the rotation classifier R_{ϕ} with the self-supervised loss L_{self} . The pipeline for relative patch location self-supervision is analog to this one.

variety of training data in this manner, one might expect to learn stronger image features and to get further performance gain in few-shot learning. At the extreme, using only unlabeled data in the first learning stage, thus removing the use of base classes altogether, is also appealing. We will show that both these semi-supervised and unsupervised regimes can be indeed put at work for few-shot object recognition thanks to self-supervised tasks.

In summary, the contributions of our work are: **(1)** We propose to weave self-supervision into the training objective of few-shot learning algorithms. The goal is to boost the ability of the latter to adapt to novel classes with few training data. **(2)** We study the impact of the added self-supervised loss by performing exhaustive quantitative experiments on Mini-Imagenet, CIFAR-FS, tiered-MiniImagenet, and ImageNet-FS few-shot datasets. In all of them, self-supervision improves the few-shot learning performance leading to state-of-the-art results. **(3)** Finally, we extend the proposed few-shot recognition framework to semi-supervised and unsupervised setups, getting further performance gain in the former, and showing with the latter that our framework provides a protocol for evaluating self-supervised representations on few-shot object recognition *which is quite simple and fast*.

2. Related work

Few-shot learning. There is a broad array of few-shot learning approaches, including, among many: gradient descent-based approaches [1, 10, 39, 45], which learn how to rapidly adapt a model to a given few-shot recognition task via a small number of gradient descent iterations; metric learning-based approaches that learn a distance metric between a query, *i.e.*, test image, and the support images, *i.e.*, training images, of a few-shot task [26, 33, 53, 55, 57, 59]; methods learning to map a test example to a class label by accessing memory modules that store training examples for

that task [11, 25, 35, 38, 50]; approaches that learn how to generate the weights of a classifier [12, 16, 43, 44, 13] or of a multi-layer neural network [3, 18, 19, 58] for the new classes given the few available training data for each of them; methods that “hallucinate” additional examples of a class from a reduced amount of data [21, 57].

In our work we consider two approaches from the metric learning category, namely *Prototypical Networks* [53] and *Cosine Classifiers* [12, 43] for their simplicity and flexibility. Nevertheless, the proposed auxiliary self-supervision is compatible with several other few-shot classification solutions.

Self-supervised learning. It is a recent paradigm for unsupervised learning that defines an annotation-free pretext task to provide a surrogate supervision signal for feature learning. Predicting the colors of images [31, 61], the relative position of image patches [6, 40], the random rotation that has been applied to an image [14], or the missing central part of an image [42], are some of the many methods [15, 32, 36, 56, 62] for self-supervised feature learning. The intuition is that, by solving such tasks, the trained model extracts semantic features that can be useful for other downstream tasks. In our case, we consider a multi-task setting where we train the backbone convnet using joint supervision from the supervised end-task and an auxiliary self-supervised pretext task. Unlike most multi-task settings aiming at good results on all tasks simultaneously [27], our goal is to improve the performance of only the main task by leveraging the supervision of the surrogate task, as also shown in [37]. We expect that, in a few-shot setting, where squeezing out generalizable features from the available data is very important, the use of self-supervision as an auxiliary task will bring improvements. Also, related to our work, Chen *et al.* [5] recently added rotation prediction self-supervision to generative adversarial networks [17] leading to significant quality improvements of the synthesized images.

3. Methodology

As already explained, few-shot learning algorithms have two learning stages and two corresponding sets of classes. Here, we define as $D_b = \{(\mathbf{x}, y)\} \subset I \times Y_b$ the training set of *base classes* used during the first learning stage, where $\mathbf{x} \in I$ is an image with label y in label set Y_b of size N_b . Also, we define as $D_n = \{(\mathbf{x}, y)\} \subset I \times Y_n$ the training set of N_n *novel classes* used during the second learning stage, where each class has K samples ($K = 1$ or 5 in benchmarks). One talks about N_n -way K -shot learning. Importantly, the label sets Y_n and Y_b are disjoint.

In the remainder of this section, we first describe in §3.1 the two standard few-shot learning methods that we consider and introduce in §3.2 the proposed method to boost their performance with self-supervision.

3.1. Explored few-shot learning methods

The main component of all few-shot algorithms is a feature extractor $F_\theta(\cdot)$, which is a convnet with parameters θ . Given an image \mathbf{x} , the feature extractor will output a d -dimensional feature $F_\theta(\mathbf{x})$. In this work we consider two representative few-shot algorithms, Prototypical Networks (PN) [53] and Cosine Classifiers (CC) [12, 43], described below. They are fairly similar, with their main difference lying in the first learning stage: only CC learns actual base classifiers along with the feature extractor, while PN simply relies on class-level averages.

Prototypical Networks (PN) [53]. During the first stage of this approach, the feature extractor $F_\theta(\cdot)$ is learned on sampled few-shot classification sub-problems that are analogue to the targeted few-shot classification problem. In each training episode of this learning stage, a subset $Y_* \subset Y_b$ of N_* base classes are sampled (they are called “support classes”) and, for each of them, K training examples are randomly picked from within D_b . This yields a training set D_* . Given current feature extractor F_θ , the average feature for each class $j \in Y_*$, its “prototype”, is computed as

$$\mathbf{p}_j = \frac{1}{K} \sum_{\mathbf{x} \in X_*^j} F_\theta(\mathbf{x}), \text{ with } X_*^j = \{\mathbf{x} \mid (\mathbf{x}, y) \in D_*, y = j\} \quad (1)$$

and used to build a simple similarity-based classifier. Then, given a new image \mathbf{x}_q from a support class but different from samples in D_* , the classifier outputs for each class j the normalized classification score

$$C^j(F_\theta(\mathbf{x}_q); D_*) = \text{softmax}_j \left[\text{sim}(F_\theta(\mathbf{x}_q), \mathbf{p}_i)_{i \in Y_*} \right], \quad (2)$$

where $\text{sim}(\cdot, \cdot)$ is a similarity function, which may be cosine similarity or negative squared Euclidean distance. So, in practice, the image \mathbf{x}_q will be classified to its closest prototype. Note that the classifier is conditioned on D_* in order to

compute the class prototypes. The first learning stage finally amounts to iteratively minimizing the following loss w.r.t. θ :

$$L_{\text{few}}(\theta; D_b) = \mathbb{E}_{\substack{D_* \sim D_b \\ (\mathbf{x}_q, y_q)}} \left[-\log C^{y_q}(F_\theta(\mathbf{x}_q); D_*) \right], \quad (3)$$

where (\mathbf{x}_q, y_q) is a training sample from a support class defined in D_* but different from images in D_* .

In the second learning stage, the feature extractor F_θ is frozen and the classifier of novel classes is simply defined as $C(\cdot; D_n)$, with prototypes defined as in (1) with $D_* = D_n$.

Cosine Classifiers (CC) [12, 43]. In CC few-shot learning, the first stage trains the feature extractor F_θ together with a cosine-similarity based classifier on the (standard) supervised task of classifying the base classes. Denoting $W_b = [\mathbf{w}_1, \dots, \mathbf{w}_{N_b}]$ the matrix of the d -dimensional classification weight vectors, the normalized score for an input image \mathbf{x} reads

$$C^j(F_\theta(\mathbf{x}); W_b) = \text{softmax}_j \left[\gamma \cos(F_\theta(\mathbf{x}), \mathbf{w}_i)_{i \in Y_b} \right], \quad (4)$$

where $\cos(\cdot, \cdot)$ is the cosine operation between two vectors, and the scalar γ is the inverse temperature parameter of the softmax operator.¹

The first learning stage aims at minimizing w.r.t. θ and W_b the negative log-likelihood loss:

$$L_{\text{few}}(\theta, W_b; D_b) = \mathbb{E}_{(\mathbf{x}, y) \sim D_b} \left[-\log C^y(F_\theta(\mathbf{x}); W_b) \right]. \quad (5)$$

One of the reasons for using the cosine-similarity based classifier instead of the standard dot-product based one, is that the former learns feature extractors that reduce intra-class variations and thus can generalize better on novel classes. By analogy with PN, the weight vectors \mathbf{w}_j ’s can be interpreted as learned prototypes for the base classes, to which input image features are compared for classification.

As with PN, the second stage boils down to computing one representative feature \mathbf{w}_j for each new class by simple averaging of associated K samples in D_n , and to define the final classifier $C(\cdot; [\mathbf{w}_1 \dots \mathbf{w}_{N_n}])$ the same way as in (4).

3.2. Boosting few-shot learning via self-supervision

A major challenge in few-shot learning is encountered during the first stage of learning. How to make the feature extractor learn image features that can be readily exploited for novel classes with few training data during the second stage? With this goal in mind, we propose to leverage the recent progress in self-supervised feature learning to further improve current few-shot learning approaches.

¹Specifically, γ controls the peakiness of the probability distribution generated by the softmax operator [23].

Through solving a non-trivial pretext task that can be trivially supervised, such as recovering the colors of images from their intensities, a network is encouraged to learn rich and generic image features that are transferable to other downstream tasks such as image classification. In the first learning stage, we propose to extend the training of the feature extractor $F_\theta(\cdot)$ by including such a self-supervised task besides the main task of recognizing base classes.

We consider two ways for incorporating self-supervision into few-shot learning algorithms: (1) by using an auxiliary loss function based on a self-supervised task, and (2) by exploiting unlabeled data in a semi-supervised way during training. We describe the two techniques in the following.

3.2.1 Auxiliary loss based on self-supervision

We incorporate self-supervision to a few-shot learning algorithm by adding an auxiliary self-supervised loss during its first learning stage. More formally, let $L_{\text{self}}(\theta, \phi; X_b)$ be the self-supervised loss applied to the set $X_b = \{\mathbf{x} \mid (\mathbf{x}, y) \in D_b\}$ of training examples in D_b deprived of their class labels. The loss $L_{\text{self}}(\theta, \phi; X_b)$ is a function of the parameters θ of the feature extractor and of the parameters ϕ of a network only dedicated to the self-supervised task. The first training stage of few-shot learning now reads

$$\min_{\theta, [W_b], \phi} L_{\text{few}}(\theta, [W_b]; D_b) + \alpha L_{\text{self}}(\theta, \phi; X_b), \quad (6)$$

where L_{few} stands either for the PN few-shot loss (3) or for the CC one (5), with additional argument W_b in the latter case (hence the bracket notation). The positive hyperparameter α controls the importance of the self-supervised term². An illustration of the approach is provided in Figure 1.

For the self-supervised loss, we consider two tasks in the present work: predicting the rotation incurred by an image, [14], which is simple and readily incorporated into a few-shot learning algorithm; predicting the relative location of two patches from the same image [6], a seminal task in self-supervised learning. In a recent study, both methods have been shown to achieve state-of-the-art results [28].

Image rotations. In this task, the convnet must recognize among four possible 2D rotations in $\mathcal{R} = \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ the one applied to an image (see Figure 1). Specifically, given an image \mathbf{x} , we first create its four rotated copies $\{\mathbf{x}^r \mid r \in \mathcal{R}\}$, where \mathbf{x}^r is the image \mathbf{x} rotated by r degrees. Based on the features $F_\theta(\mathbf{x}^r)$ extracted from such a rotated image, the new network R_ϕ attempts to predict the rotation class r . Accordingly, the self-supervised loss of this task is defined as:

$$L_{\text{self}}(\theta, \phi; X) = \mathbb{E}_{\mathbf{x} \sim X} \left[\sum_{\forall r \in \mathcal{R}} -\log R_\phi^r(F_\theta(\mathbf{x}^r)) \right], \quad (7)$$

²In our experiments, we use $\alpha = 1.0$.

where X is the original training set of non-rotated images and $R_\phi^r(\cdot)$ is the predicted normalized score for rotation r . Intuitively, in order to do well for this task the model should reduce the bias towards *up-right* oriented images, typical for ImageNet-like datasets, and learn more diverse features to disentangle classes in the low-data regime.

Relative patch location. Here, we create random pairs of patches from an image and then predict, among eight possible positions, the location of the second patch w.r.t. to the first, *e.g.*, “on the left and above” or “on the right and below”. Specifically, given an image \mathbf{x} , we first divide it into 9 regions over a 3×3 grid and sample a patch within each region. Let’s denote $\bar{\mathbf{x}}^0$ the central image patch, and $\bar{\mathbf{x}}^1 \dots \bar{\mathbf{x}}^8$ its eight neighbors lexicographically ordered. We compute the representation of each patch³ and then generate patch feature pairs $(F_\theta(\bar{\mathbf{x}}^0), F_\theta(\bar{\mathbf{x}}^p))$ by concatenation. We train a fully-connected network $P_\phi(\cdot, \cdot)$ to predict the position of $\bar{\mathbf{x}}^p$ from each pair.

The self-supervised loss of this task is defined as:

$$L_{\text{self}}(\theta, \phi; X) = \mathbb{E}_{\mathbf{x} \sim X} \left[\sum_{p=1}^8 -\log P_\phi^p(F_\theta(\bar{\mathbf{x}}^0), F_\theta(\bar{\mathbf{x}}^p)) \right], \quad (8)$$

where X is a set of images and P_ϕ^p is the predicted normalized score for the relative location p . Intuitively, a good model on this task should somewhat recognize objects and parts, even in presence of occlusions and background clutter. Note that, in order to prevent the model from learning low-level image statistics such as chromatic aberration [6], the patches are preprocessed with aggressive color augmentation (*i.e.*, converting to grayscale with probability 0.66 and normalizing the pixels of each patch individually to have zero mean and unit standard deviation).

3.2.2 Semi-supervised few-shot learning

The self-supervised term L_{self} in the training loss (6) does not depend on class labels. We can easily extend it to learn as well from additional unlabeled data. Indeed, if a set X_u of unlabeled images is available besides D_b , we can make the self-supervised task benefit from them by redefining the first learning stage as:

$$\min_{\theta, [W_b], \phi} L_{\text{few}}(\theta, [W_b]; D_b) + \alpha L_{\text{self}}(\theta, \phi; X_b \cup X_u). \quad (9)$$

By training the feature extractor F_θ to also minimize the self-supervised loss on these extra unlabeled images, we open up its visual scope with the hope that this will further improve its ability to accommodate novel classes with scarce data. This can be seen as a semi-supervised training approach

³If the architecture of $F_\theta(\cdot)$ is fully convolutional, we can apply it to both big images and smaller patches.

for few-shot algorithms. An interesting aspect of this semi-supervised training approach is that it does not require the extra unlabeled data to be from the same (base) classes as those in labeled dataset D_b . Thus, it is much more flexible w.r.t. the source of the unlabeled data than standard semi-supervised approaches.

4. Experimental Results

In this section we evaluate self-supervision as auxiliary loss function in §4.2 and in §4.3 as a way of exploiting unlabeled data in semi-supervised training. Finally, in §4.4 we use the few-shot object recognition task for evaluating self-supervised methods. We make available more experimental results along with implementation details in the Appendices A and B provided as supplementary material.

Datasets. We perform experiments on four few-shot datasets, MiniImageNet [55], *tiered*-MiniImageNet [46], CIFAR-FS [2], and ImageNet-FS [20]. **MiniImageNet** consists of 100 classes randomly picked from the ImageNet dataset [48] (*i.e.*, 64 base classes, 16 validation classes, and 20 novel test classes); each class has 600 images with size 84×84 pixels. **tiered-MiniImageNet** consists of 608 classes randomly picked from ImageNet (*i.e.*, 351 base classes, 97 validation classes, and 160 novel test classes); in total there are 779, 165 images again with size 84×84 . Finally, **ImageNet-FS** is a few-shot benchmark based on the original ImageNet with 389 base classes, 300 novel validation classes, and 311 novel test classes. **CIFAR-FS** is a few-shot dataset created by dividing the 100 classes of CIFAR-100 into 64 base classes, 16 validation classes, and 20 novel test classes. The images in this dataset have size 32×32 pixels.

Evaluation metrics. Few-shot classification algorithms are evaluated based on the classification accuracy in their second learning stage (when the learned classifier is applied to test images from the novel classes). More specifically, a large number of N_n -way K -shot tasks are sampled from the available set of novel classes. Each task is created by randomly selecting N_n novel classes from the available test (validation) classes and then within the selected classes randomly selecting K train and M test images per class (making sure that train and test images do not overlap). The classification accuracy is measured on the $N_n \times M$ test images and is averaged over all the sampled few-shot tasks. Except otherwise stated, for all experiments we used $M = 15$, $N_n = 5$, and $K = 1$ or $K = 5$ (1-shot and 5-shot settings respectively). ImageNet-FS is a more challenging benchmark: it evaluates a method on the validation set of ImageNet based on the classification accuracy of both novel and base classes using all the $N_n = 311$ novel classes at once.

4.1. Implementation details

Network architectures. We use four different feature extractor architectures F_θ : Conv-4-64 [55], Conv-4-512, WRN-28-10 [60], and ResNet10 [22]. The network $R_\phi(\cdot)$ specific to the rotation prediction task gets as input the output feature maps of F_θ and is implemented as a convnet. Given two patches, the network $P_\phi(\cdot, \cdot)$ specific to the relative patch location task gets the concatenation of their feature vectors extracted with F_θ as input, and forwards it to two fully connected layers. For further details see Appendix B.1.

Training optimization routine for first learning stage.

We optimize the training loss with mini-batch stochastic gradient descent (SGD). For the labeled data we apply both recognition L_{few} and self-supervised L_{self} losses. For the semi-supervised training, at each step we sample mini-batches that consist of labeled data, for which we use both losses, and unlabeled data, in which case we apply only L_{self} . For further details see Appendices B.2 and B.3.

Implementation of relative patch location task.

Due to the aggressive color augmentation of the patches in the patch localization task, and the fact that the patches are around 9 times smaller than the original images, the data distribution that the feature extractor “sees” from them is very different from that of the images. To overcome this problem we apply an extra auxiliary classification loss to the features extracted from the patches. Specifically, during the first learning stage of CC we merge the features $F_\theta(\bar{x}^p)$ of the 9 patches of an image (*e.g.*, with concatenation or averaging) and then apply the cosine classifier (4) to the resulting feature (this classifier does not share its weight vectors with the classifier applied to the original images features). Note that this patch based auxiliary classification loss has the same weight as the original classification loss L_{few} . Also, during the second learning stage we do not use the patch based classifier.

4.2. Self-supervision as auxiliary loss function

Rotation prediction as auxiliary loss. We first study the impact of adding rotation prediction as self-supervision to the few-shot learning algorithms of Cosine Classifiers (CC) and Prototypical Networks (PN). We perform this study using the MiniImageNet and CIFAR-FS datasets and report results in Tables 1 and 2 respectively. For the CC case, we use as strong baselines, CC models without self-supervision but trained to recognize all the 4 rotated versions of an image. The reason for using this baseline is that during the first learning stage, the model “sees” the same type of data, *i.e.*, rotated images, as the model with rotation prediction self-supervision. Note that despite the rotation augmentations of the first learning stage, during the second stage the model *uses as training examples for the novel classes only*

Model	Backbone	1-shot	5-shot
CC	Conv-4-64	54.31 ± 0.42%	70.89 ± 0.34%
CC+rot		54.83 ± 0.43%	71.86 ± 0.33%
CC	Conv-4-512	55.68 ± 0.43%	73.19 ± 0.33%
CC+rot		56.27 ± 0.43%	74.30 ± 0.33%
CC	WRN-28-10	61.09 ± 0.44%	78.43 ± 0.33%
CC+rot		62.93 ± 0.45%	79.87 ± 0.33%
PN	Conv-4-64	52.20 ± 0.46%	69.98 ± 0.36%
PN+rot		53.63 ± 0.43%	71.70 ± 0.36%
PN	Conv-4-512	54.60 ± 0.46%	71.59 ± 0.36%
PN+rot		56.02 ± 0.46%	74.00 ± 0.35%
PN	WRN-28-10	55.85 ± 0.48%	68.72 ± 0.36%
PN+rot		58.28 ± 0.49%	72.13 ± 0.38%

Table 1: Rotation prediction as auxiliary loss on MiniImageNet. Average 5-way classification accuracies for the novel classes with 95% confidence intervals (using 2000 episodes).

Models	Backbone	1-shot	5-shot
CC	Conv-4-64	61.80 ± 0.30%	78.02 ± 0.24%
CC+rot		63.45 ± 0.31%	79.79 ± 0.24%
CC	Conv-4-512	65.26 ± 0.31%	81.14 ± 0.23%
CC+rot		65.87 ± 0.30%	81.92 ± 0.23%
CC	WRN-28-10	74.51 ± 0.31%	86.45 ± 0.22%
CC+rot		75.38 ± 0.31%	87.25 ± 0.21%
PN	Conv-4-64	62.82 ± 0.32%	79.59 ± 0.24%
PN+rot		64.69 ± 0.32%	80.82 ± 0.24%
PN	Conv-4-512	66.48 ± 0.32%	80.28 ± 0.23%
PN+rot		67.94 ± 0.31%	82.20 ± 0.23%
PN	WRN-28-10	69.48 ± 0.34%	82.43 ± 0.23%
PN+rot		69.55 ± 0.34%	82.34 ± 0.24%

Table 2: Rotation prediction as auxiliary loss on CIFAR-FS. Average 5-way classification accuracies for the novel classes with 95% confidence intervals (using 5000 episodes).

the up-right versions of the images. Still however, using rotation augmentations improves the classification performance of the baseline models when adapted to the novel classes. Therefore, for fair comparison, we also apply rotation augmentations to the CC models with rotation prediction self-supervision. For the PN case, we do not use rotation augmentation since in our experiments this lead to performance degradation.

The results in Tables 1 and 2 demonstrate that (1) indeed, adding rotation prediction self-supervision improves the few-shot classification performance, and (2) the performance improvement is more significant for high capacity architectures, e.g., WRN-28-10. In Table 3 we provide results for the challenging ImageNet-FS [20] benchmark, using a ResNet10 backbone and CC. Our method achieves distinct and consistent improvements on the recognition of novel

Model	Classes	1-shot	2-shot	5-shot	10-shot	20-shot
CC	Novel	45.00 ± 0.24	56.33 ± 0.17	68.13 ± 0.09	73.24 ± 0.07	76.03 ± 0.05
CC+rot		46.43 ± 0.24	57.80 ± 0.16	69.67 ± 0.09	74.64 ± 0.06	77.31 ± 0.05
CC	All	56.87 ± 0.15	63.66 ± 0.10	71.18 ± 0.06	74.51 ± 0.04	76.10 ± 0.03
CC+rot		57.88 ± 0.15	64.76 ± 0.10	72.29 ± 0.07	75.63 ± 0.04	77.40 ± 0.03

Table 3: Rotation prediction as auxiliary loss on ImageNet-FS [20]. Top-5 accuracy results with ResNet10 backbone. The accuracy on base classes is 92.91% for CC and 93.52% for CC+rot. We did not use rotation augmentations for these results.

Model	Backbone	1-shot	5-shot
CC	Conv-4-64	53.72 ± 0.42%	70.96 ± 0.33%
CC+loc		54.30 ± 0.42%	71.58 ± 0.33%
CC	Conv-4-512	55.58 ± 0.42%	73.52 ± 0.33%
CC+loc		56.87 ± 0.42%	74.84 ± 0.33%
CC	WRN-28-10	58.43 ± 0.46%	75.45 ± 0.34%
CC+loc		60.71 ± 0.46%	77.64 ± 0.34%
PN	WRN-28-10	55.88 ± 0.49%	70.12 ± 0.40%
PN+loc		58.16 ± 0.49%	72.89 ± 0.39%

Table 4: Relative patch location as auxiliary loss on MiniImageNet. Average 5-way classification accuracies for the novel classes with 95% confidence intervals (using 2000 episodes).

classes alone as well as on the simultaneous recognition of both novel and base classes (All).

Relative patch location prediction as auxiliary loss. As explained in §3.2.1, we consider a second self-supervised task, namely relative patch location prediction. For simplicity, we restrict its assessment to the CC algorithm for the Conv-4-64 and Conv-4-512 backbones, which in our experiments proved to perform better than PN and to be simpler to train. Also, we consider only the MiniImageNet dataset and not CIFAR-FS since the latter contains thumbnail images of size 32×32 from which it does not make sense to extract patches: their size would have to be less than 8×8 pixels, which is too small for the evaluated architectures. We report results on MiniImageNet in Table 4. As a strong baseline we used CC models without self-supervision but with the auxiliary patch-based classification loss described in §4.1.

Based on the results of Table 4 we observe that: (1) relative patch location also manages to improve the few-shot classification performance and, as in the rotation prediction case, the improvement is more significant for high capacity network architectures. (2) Also, comparing to the rotation prediction case, the relative patch location, in general, offers smaller performance improvement.

Comparison with prior work. In Tables 5, 6, and 7, we compare our approach with prior few-shot methods on the MiniImageNet, CIFAR-FS, and *tiered*-MiniImageNet datasets respectively. For our approach we used CC and rotation prediction self-supervision, which before gave the

Models	Backbone	1-shot	5-shot
MAML [10]	Conv-4-64	48.70 ± 1.84%	63.10 ± 0.92%
Prototypical Nets [53]	Conv-4-64	49.42 ± 0.78%	68.20 ± 0.66%
Lwof [12]	Conv-4-64	56.20 ± 0.86%	72.81 ± 0.62%
RelationNet [59]	Conv-4-64	50.40 ± 0.80%	65.30 ± 0.70%
GNN [11]	Conv-4-64	50.30%	66.40%
R2-D2 [2]	Conv-4-64	48.70 ± 0.60%	65.50 ± 0.60%
R2-D2 [2]	Conv-4-512	51.20 ± 0.60%	68.20 ± 0.60%
TADAM [41]	ResNet-12	58.50 ± 0.30%	76.70 ± 0.30%
Munkhdalai <i>et al.</i> [38]	ResNet-12	57.10 ± 0.70%	70.04 ± 0.63%
SNAIL [51]	ResNet-12	55.71 ± 0.99%	68.88 ± 0.92%
Qiao <i>et al.</i> [44]*	WRN-28-10	59.60 ± 0.41%	73.74 ± 0.19%
LEO [49]*	WRN-28-10	61.76 ± 0.08%	77.59 ± 0.12%
CC+rot	Conv-4-64	54.83 ± 0.43%	71.86 ± 0.33%
CC+rot	Conv-4-512	56.27 ± 0.43%	74.30 ± 0.34%
CC+rot	WRN-28-10	62.93 ± 0.45%	79.87 ± 0.33%
CC+rot+unlabeled	WRN-28-10	64.03 ± 0.46%	80.68 ± 0.33%

Table 5: Comparison with prior work on MiniImageNet. *: using also the validation classes for training. For the description of the CC+rot+unlabeled model see §4.3.

Models	Backbone	1-shot	5-shot
PN [53]‡	Conv-4-64	62.82 ± 0.32%	79.59 ± 0.24%
PN [53]‡	Conv-4-512	66.48 ± 0.32%	80.28 ± 0.23%
MAML [10]†	Conv-4-64	58.90 ± 1.90%	71.50 ± 1.00%
MAML [10]†	Conv-4-512	53.80 ± 1.80%	67.60 ± 1.00%
RelationNet [59]†	Conv-4-64	55.00 ± 1.00%	69.30 ± 0.80%
GNN [11]†	Conv-4-64	61.90%	75.30%
GNN [11]†	Conv-4-512	56.00%	72.50%
R2-D2 [2]	Conv-4-64	60.00 ± 0.70%	76.10 ± 0.60%
R2-D2 [2]	Conv-4-512	64.00 ± 0.80%	78.90 ± 0.60%
CC+rot	Conv-4-64	63.45 ± 0.31%	79.79 ± 0.24%
CC+rot	Conv-4-512	65.87 ± 0.30%	81.92 ± 0.23%
CC+rot	WRN-28-10	76.09 ± 0.30%	87.83 ± 0.21%

Table 6: Comparison with prior work in CIFAR-FS. †: results from [2]. ‡: our implementation.

Models	Backbone	1-shot	5-shot
MAML [10]†	Conv-4-64	51.67 ± 1.81%	70.30 ± 0.08%
Prototypical Nets [53]	Conv-4-64	53.31 ± 0.89%	72.69 ± 0.74%
RelationNet [59]†	Conv-4-64	54.48 ± 0.93%	71.32 ± 0.78%
Liu <i>et al.</i> [34]	Conv-4-64	57.41 ± 0.94%	71.55 ± 0.74%
LEO [49]	WRN-28-10	66.33 ± 0.05%	81.44 ± 0.09%
CC	WRN-28-10	70.04 ± 0.51%	84.14 ± 0.37%
CC+rot	WRN-28-10	70.53 ± 0.51%	84.98 ± 0.36%

Table 7: Rotation prediction as auxiliary loss on tiered-MiniImageNet. Average 5-way accuracies for the novel classes with 95% confidence (using 2000 episodes). †: results from [34].

best results. In all cases we achieve state-of-the-art results surpassing prior methods with a significant margin. For instance, in the 1-shot and 5-shot settings of MiniImageNet we outperform the previous leading method LEO [49] by

Model	MiniImageNet		CIFAR-FS	
	1-shot	5-shot	1-shot	5-shot
CC+rot	62.93 ± 0.46%	79.87 ± 0.33%	75.38 ± 0.31%	87.25 ± 0.21%
rot→CC	61.95 ± 0.45%	79.25 ± 0.35%	74.90 ± 0.31%	86.72 ± 0.22%

Table 8: Self-supervised (rotation) pretraining with WRN-28-10.

around 1.2 and 2.3 percentage points respectively.

Comparison with self-supervised pre-training. An alternative way to exploit self-supervision in few-shot learning is by self-supervised pre-training of the feature extractor of a few-shot model. In Table 8 we provide such results for CC and the rotation prediction task (entry ‘rot→CC’) using the high-capacity WRN-28-10 backbone. We observe that it offers smaller improvements than the advocated approach.

More detailed results are provided in Appendix A.

4.3. Semi-supervised few-shot learning

Next, we evaluate the proposed semi-supervised training approach. In these experiments we use CC models with rotation prediction self-supervision. We perform two types of semi-supervised experiments: (1) training with unlabeled data from the same base classes, and (2) training with unlabeled data that are not from the base classes.

Training with unlabeled data from the same base classes.

From the base classes of MiniImageNet, we use only a small percentage of the training images (*e.g.*, 5% of images per class) as annotated training data, while the rest of the images (*e.g.*, 95%) are used as the unlabeled data in the semi-supervised training. We provide results in the first two sections of Table 9. The proposed semi-supervised training approach is compared with a CC model without self-supervision and with a CC model with self-supervision but no recourse to the unlabeled data. The results demonstrate that indeed, our method manages to improve few-shot classification performance by exploiting unlabeled images. Compared to Ren *et al.* [46], which also propose a semi-supervised method, our method with Conv-4-64 and 20% annotations achieves better results than their method with Conv-4-64 and 40% annotations (*i.e.*, our **51.21%** and **68.89%** MiniImageNet accuracies vs. their 50.41% and 64.39% for the 1-shot and 5-shot settings respectively).

Training with unlabeled data not from the base classes.

This is a more realistic setting, since it is hard to constrain the unlabeled images to be from the same classes as the base classes. For this experiment, we used as unlabeled data the training images of the *tiered*-MiniImageNet base classes *minus the classes that are common with the base, validation, or test classes of MiniImageNet*. In total, 408, 726 unlabeled images are used from *tiered*-MiniImageNet. We

Rot	M	T	$\mu = 5\%$		$\mu = 10\%$		$\mu = 20\%$	
			1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Conv-4-64			41.87%	57.76%	45.63%	62.29%	49.34%	66.48%
✓			43.26%	58.88%	47.57%	64.03%	50.48%	67.92%
✓	✓		45.41%	62.14%	47.75%	64.93%	51.21%	68.89%
WRN-28-10			39.05%	54.89%	44.59%	61.42%	49.52%	67.81%
✓			43.60%	60.52%	49.05%	67.35%	53.66%	72.69%
✓	✓		47.25%	65.07%	52.90%	71.48%	56.79%	74.67%
✓		✓	46.95%	64.93%	52.66%	71.13%	55.37%	73.66%

Table 9: Semi-supervised training with rotation prediction as self-supervision on MiniImageNet. μ is the percentage of the base class training images of MiniImageNet that are used as annotated data during training. *Rot* indicates adding self-supervision, *M* indicates using as unlabeled data the (rest of) MiniImageNet training dataset, and *T* using as unlabeled data the *tiered*-MiniImageNet training dataset.

Models	Backbone	1-shot	5-shot	20-shot	50-shot
CACTUs [24]	Conv-4-64	39.90%	53.97%	63.84%	69.64%
CC	Conv-4-64	53.63%	70.74%	80.03%	82.61%
Rot		41.70%	58.64%	68.61%	71.86%
Loc		37.75%	53.02%	61.38%	64.15%
CC	WRN-28-10	58.59%	76.59%	82.70%	84.27%
Rot		43.43%	60.86%	69.82%	72.20%
Loc		41.78%	59.10%	67.86%	70.32%

Table 10: Evaluating self-supervised representation learning methods on few-shot recognition on MiniImageNet. *Rot* refers to the rotation prediction task, *Loc* to the relative patch location task, and *CC* to the supervised method of Cosine Classifiers.

report results in the last row of Table 9. Indeed, even in this difficult setting, our semi-supervised approach is still able to exploit unlabeled data and to improve the classification performance. Furthermore, we did an extra experiment in which we trained a WRN-28-10 based model using 100% of MiniImageNet training images and unlabeled data from *tiered*-MiniImageNet. This model achieved **64.03%** and **80.68%** accuracies for the 1-shot and 5-shot settings respectively on MiniImageNet (see entry ‘CC+rot+unlabeled’ of Table 5), which improves over the already very strong CC+rot model (see Table 5).

4.4. Few-shot object recognition to assess self-supervised representations

Given that our framework allows the easy combination of any type of self-supervised learning approach with the adopted few-shot learning algorithms, we also propose to use it as an alternative way for comparing/evaluating the effectiveness of different self-supervised approaches. To this end, the only required change to our framework is to use

uniquely the self-learning loss (*i.e.*, no labeled data is now used) in the first learning stage (for implementation details see Appendix B.2). The performance of the few-shot model resulting from the second learning stage can then be used for evaluating the self-supervised method under consideration.

Comparing competing self-supervised techniques is not straightforward since it must be done by setting up another, somewhat contrived task that exploits the learned representations [6, 28]. Instead, given the very similar goals of few-shot and self-supervised learning, we argue that the proposed comparison method could be more meaningful for assessing different self-supervised features. Furthermore, (1) *it obtains more robust statistics* aggregated over evaluations of thousands of episodes with multiple different configurations of classes and training/testing samples, (2) *it is very fast* (just a few minutes to evaluate), and (3) *it is significantly simpler than other protocols* (no need to adapt the architecture to a new task or fine-tune parameters, hence no hyper-parameter tuning). Those advantages might make the protocol practically useful to self-supervision research.

To illustrate our point, we provide in Table 10 quantitative results of this type of evaluation on the MiniImageNet dataset, for the self-supervision methods of rotation prediction and relative patch location prediction. For self-supervised training we used the training images of the base classes of MiniImageNet and for the few-shot classification step we used the test classes of MiniImageNet. We observe that the explored self-supervised approaches achieve relatively competitive classification performance when compared to the supervised method of CC and obtain results that are on par or better than other, more complex, unsupervised systems. We leave as future work a more detailed and thorough comparison of self-learned representations in this evaluation setting.

5. Conclusions

Inspired by the close connection between few-shot and self-supervised learning, we propose to add an auxiliary loss based on self-supervision during the training of few-shot recognition models. The goal is to boost the ability of the latter to recognize novel classes using only few training data. Our detailed experiments on MiniImageNet, CIFAR-FS, *tiered*-MiniImageNet, and ImageNet-FS few-shot datasets reveal that indeed adding self-supervision leads to significant improvements on the few-shot classification performance, which makes the employed few-shot models achieve state-of-the-art results. Furthermore, the annotation-free nature of the self-supervised loss allows us to exploit diverse unlabeled data in a semi-supervised manner, which further improves the classification performance. Finally, we show that the proposed framework can also be used for evaluating self-supervised or unsupervised methods based on few-shot object recognition.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *NIPS*, 2016. 2
- [2] Luca Bertinetto, João Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR*, 2019. 5, 7
- [3] Luca Bertinetto, João Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *NIPS*, 2016. 2
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans. PAMI*, 40(4):834–848, 2018. 1
- [5] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised generative adversarial networks. *arXiv preprint arXiv:1811.11212*, 2018. 2
- [6] Carl Doersch, Abhinav Gupta, and Alexei Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 1, 2, 4, 8
- [7] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, pages 766–774, 2014. 1
- [8] Li Fei-Fei, Rob Fergus, and Pietro Perona. A Bayesian approach to unsupervised one-shot learning of object categories. In *ICCV*, 2003. 1
- [9] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Trans. PAMI*, 28(4):594–611, 2006. 1
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. 2, 7
- [11] Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*, 2017. 2, 7
- [12] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 2, 3, 7
- [13] Spyros Gidaris and Nikos Komodakis. Generating classification weights with GNN denoising autoencoders for few-shot learning. In *CVPR*, 2019. 2
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 1, 2, 4
- [15] Clément Godard, Oisín Mac Aodha, and Gabriel Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 2
- [16] Faustino Gomez and Jürgen Schmidhuber. Evolving modular fast-weight networks for control. In *ICANN*, 2005. 2
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [18] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. 2
- [19] Chunrui Han, Shiguang Shan, Meina Kan, Shuzhe Wu, and Xilin Chen. Face recognition with contrastive convolution. In *ECCV*, 2018. 2
- [20] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. *arXiv preprint arXiv:1606.02819*, 2016. 5, 6
- [21] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, 2017. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 5
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [24] Kyle Hsu, Sergey Levine, and Chelsea Finn. Unsupervised learning via meta-learning. In *ICLR*, 2019. 8
- [25] Łukasz Kaiser, Ofir Nachum, Aurko Roy, and Samy Bengio. Learning to remember rare events. *arXiv preprint arXiv:1703.03129*, 2017. 2
- [26] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, 2015. 1, 2
- [27] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *CVPR*, 2017. 2
- [28] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. Revisiting self-supervised visual representation learning. *arXiv preprint arXiv:1901.09005*, 2019. 4, 8
- [29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [30] Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. One shot learning of simple visual concepts. In *An. Meeting of the Cognitive Science Society*, 2011. 1
- [31] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016. 1, 2
- [32] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017. 2
- [33] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. Dense classification and implanting for few-shot learning. In *CVPR*, 2019. 2
- [34] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, and Yi Yang. Transductive propagation network for few-shot learning. *arXiv preprint arXiv:1805.10002*, 2018. 7
- [35] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *ICLR*, 2018. 2
- [36] Ishan Misra, Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. 2
- [37] Taylor Mordan, Nicolas Thome, Gilles Henaff, and Matthieu Cord. Revisiting multi-task learning with rock: a deep residual auxiliary block for visual detection. In *NIPS*, 2018. 2

- [38] Tsendsuren Munkhdalai and Hong Yu. Meta networks. *arXiv preprint arXiv:1703.00837*, 2017. 2, 7
- [39] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *CoRR*, abs/1803.02999, 2018. 2
- [40] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 1, 2
- [41] Boris N Oreshkin, Alexandre Lacoste, and Pau Rodriguez. Tadam: Task dependent adaptive metric for improved few-shot learning. *arXiv preprint arXiv:1805.10123*, 2018. 7
- [42] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [43] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *CVPR*, 2018. 2, 3
- [44] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan Yuille. Few-shot image recognition by predicting parameters from activations. *arXiv preprint arXiv:1706.03466*, 2, 2017. 2, 7
- [45] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. *ICLR*, 2017. 2
- [46] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua Tenenbaum, Hugo Larochelle, and Richard Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018. 5, 7
- [47] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1
- [48] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 5
- [49] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018. 7
- [50] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016. 2
- [51] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. A simple neural network module for relational reasoning. In *NIPS*, 2017. 7
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [53] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. 2, 3, 7
- [54] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 1
- [55] Oriol Vinyals, Charles Blundell, Tim Lillicrap, and Daan Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. 1, 2, 5
- [56] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *ECCV*, 2018. 2
- [57] Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. *arXiv preprint arXiv:1801.05401*, 2018. 2
- [58] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *NIPS*, 2017. 2
- [59] Flood Sung Yongxin Yang, Li Zhang, Tao Xiang, Philip Torr, and Timothy Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018. 2, 7
- [60] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 5
- [61] Richard Zhang, Phillip Isola, and Alexei Efros. Colorful image colorization. In *ECCV*, 2016. 1, 2
- [62] Tinghui Zhou, Matthew Brown, Noah Snavely, and David Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 2